arXiv:1406.7808v3 [math.NA] 11 Aug 2015

SEGMENTAL REFINEMENT: A MULTIGRID TECHNIQUE FOR DATA LOCALITY

MARK F. ADAMS, JED BROWN, MATT KNEPLEY, AND RAVI SAMTANEY

Abstract. We investigate a domain decomposed multigrid technique, segmental refinement, for solving general nonlinear elliptic boundary value problems. Brandt and Diskin first proposed this method in 1994; we continue this work by analytically and experimentally investigating its complexity. We confirm that communication of traditional parallel multigrid can be eliminated on fine grids with modest amounts of extra work and storage while maintaining the asymptotic exactness of full multigrid, although we observe a dependence on an additional parameter not considered in the original analysis. We present a communication complexity analysis that quantifies the communication costs ameliorated by segmental refinement and report performance results with up to 64K cores of a Cray XC30.

1. Introduction

Full multigrid (FMG) is a provably asymptotically exact, non-iterative algebraic equation solver for discretized elliptic partial differential equations (PDEs) with work complexity of about five residual calculations, or what is known as textbook multigrid efficiency, for the constant coefficient Laplacian [BD81]. While textbook multigrid efficiency is only provable for a few classes of elliptic problems, it has been observed experimentally in many more problems [TDB01, TOS01, ASB10], and is applicable to general nonlinear elliptic equations. Multigrid methods are widely used in practice; they are important methods to adapt to emerging architectures.

Memory movement, in both intra-node and inter-node communication, and global data dependencies are the primary drivers of costs, in power and time, for PDE simulations on current, and anticipated future, computer architectures. Memory movement pressures are not new and have been accumulating for decades, but the recent prominence of energy costs in powering memory and moving data is exacerbating this problem. Segmental refinement addresses the challenges posed by the deep memory hierarchies of modern architectures at a fundamental, algorithmic level by exploiting the local nature of multigrid processes and a tolerance for finite algebraic error, which nonetheless vanishes asymptotically. A segmental refinement data model or method explicitly decouples subdomain processing, at some level of the memory and multigrid hierarchy, which improves data locality, amortizes latency costs, and reduces data dependencies.

Brandt proposed segmental refinement in the 1970s [Bra77] §7.5; [BL11] §8.7; [Din79] as a low memory complexity technique for FMG that does not store the entire solution in

Key words and phrases. multigrid,parallel multigrid,distributed memory multigrid,segmental refinement.

Scaleable Solvers Group, Lawrence Berkeley National Laboratory, Berkeley, CA.

Mathematics and Computer Science Division, Argonne National Laboratory .

Computation Institute, University of Chicago.

King Abdullah University of Science and Technology.

memory at any one time. Brandt and Diskin recognized that segmental refinement has attractive properties for distributed memory computing [BD94]; it is inherently asynchronous and highly parallel, with no interprocess communication on the finest grids, and it requires only modest amounts of extra storage and work in buffer cells. This paper continues the development of segmental refinement by quantifying its complexity both experimentally and analytically. We present the first published multilevel numerical results, and report preliminary performance results, of a cell centered segmental refinement data model.

2. Differential and discretized problems

We consider general nonlinear elliptic problems in an open domain Ω with boundary $\partial\Omega$ of the form

$$Lu\,(x) = f\,(x) \qquad (x \in \Omega)\,, \tag{1}$$

where f is a known function, u is unknown, and L is a uniformly elliptic operator. While the methods described herein are generally applicable, we restrict ourselves to the 3D Poisson operator: $x = (x_1, x_2, x_3)$, $L = \left(\frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2} + \frac{\partial^2 u}{\partial x_3^2}\right)$. In addition to this interior equation, a suitable boundary condition on $\partial\Omega$ is assumed; we assume $u(x) = 0$; $x \in \partial\Omega$.

The discretization of equation (1) can be fairly general but we restrict ourselves to a cell-centered finite difference method on isotropic Cartesian grids and rectangular domains. For the grid Ω_h with mesh spacing h covering the domain Ω, the equation can be written as

$$L_h u_h(i) = f_h(i) \qquad (i \in \Omega_h)\,, \tag{2}$$

where $i = \left(x - \frac{h}{2}\right)/h$ is an integer vector, $x = ih + \frac{h}{2}$ is a cell center, and the boundary $\partial\Omega$ lines up with the cell edges. In 3D $i = (i_1, i_2, i_3)$ is an index for a cell in grid Ω_h. The indexing in equation (2) is dropped and field variables (e.g., u_h) are vectors of scalars.

Our grids Ω_h and subsequent subdomains are isotropic and can for the most part be expressed as tensor products of 1D grids. The lengths of Ω_h, in each dimension, is an integer vector; we simplify the presentation by using the integer N, because we use cubical subdomains. Multigrid utilizes an accurate and inexpensive solver on Ω_0 and a sequence of grids Ω_0, Ω_1, Ω_2,..., Ω_M, where $\Omega_k \equiv \Omega_{h_k}$, $h_k = h_{k-1}/2$, $h_M = h$, $N_k = 2N_{k-1}$, $N_M = N$.

3. Multigrid background

The antecedents of modern multigrid go back at least to Southwell in the 1930s [Sou40], Fedorenko in the early 1960s [Fed61], and others [TOS01]. Brandt developed multigrid's modern form in the 1970s, an asymptotically exact solver with work complexity of a few residual calculations – what is known as textbook multigrid efficiency. He applied multigrid to complex domains, variable coefficients, and nonlinear problems [Bra73]. A substantial body of literature, both theoretical and experimental, demonstrates the efficacy of multigrid [TOS01, BL11]. Full Approximation Scheme (or Storage, FAS) multigrid has also been demonstrated to be an effective nonlinear solver, with costs similar to those of a linearized multigrid solve (e.g., [TOS01] §5.3.3, [ASB10]).

3.1. **Multigrid algorithm.** Multigrid starts with the observation that errors that are poorly resolved with local processes can often be resolved with local processes on a lower resolution discretization. This lower resolution problem is known as a coarse grid. Multigrid, by definition, applies this process recursively until the problem size is small enough to be solved inexpensively and exactly. The coarse grid space is represented algebraically by the columns of the prolongation operator I_H^h or I_{k-1}^k, where h is the fine grid mesh spacing and H is the coarse grid mesh spacing. Residuals are mapped from the fine grid to the coarse grid with the restriction operator I_h^H. The coarse grid operator can be formed in one of two ways (with some exceptions), either algebraically to form Galerkin (or variational) coarse grids, $L_H = I_h^H L_h I_H^h$, or by creating a new operator on each coarse grid if an explicit coarse grid with boundary conditions is available.

Correction Scheme (CS) multigrid, where coarse grids compute corrections to the solution, is appropriate for linear problems, but FAS multigrid is more natural for segmental refinement. FAS is derived by writing the coarse grid residual equation for equation (2) as

$$r_H = L_H(u_H) - L_H(\hat{u}_H) = L_H(\hat{u}_H + e_H) - L_H(\hat{u}_H), \tag{3}$$

where u_H is the exact solution, $\hat{u}_H$ is an approximation to $I_h^H u_h$ (which is the full solution represented on the coarse grid), and e is the error. With an approximate solution on the fine grid $\tilde{u}_h$, the coarse grid equation can be written as

$$L_H\left(I_h^H \tilde{u}_h + e_H\right) = L_H\left(I_h^H \tilde{u}_h\right) + I_h^H\left(f_h - L_h \tilde{u}_h\right) = f_H = I_h^H\left(f_h\right) + \tau_h^H, \tag{4}$$

and is solved approximately; τ_h^H is the tau correction, which represents a correction to the coarse grid from the fine grid. After $I_h^H \tilde{u}_h$ is subtracted from the $I_h^H \tilde{u}_h + e_H$ term the correction is applied to the fine grid with the standard prolongation process. Figure 1 shows an FAS multigrid $V(\nu 1, \nu 2)$–cycle algorithm with nonlinear local process or smoother $u \leftarrow S(L, u, f)$.

```
u = function FASMGV(L_k, u_k, f_k)
    if k > 0
        u_k ← S^ν1(L_k, u_k, f_k)
        r_k ← f_k − L_k u_k
        u_{k−1} ← Î_k^{k−1}(u_k)
        r_{k−1} ← I_k^{k−1}(r_k)
        t_{k−1} ← u_{k−1}
        w_{k−1} ← FASMGV(L_{k−1}, u_{k−1}, r_{k−1} + L_{k−1} u_{k−1})
        u_k ← u_k + I_{k−1}^k(w_{k−1} − t_{k−1})
        u_k ← S^ν2(L_k, u_k, f_k)
    else
        u_k ← L_k^{−1} f_k
    return u_k
```

FIGURE 1. FAS multigrid *V-cycle*

A lower order restriction operator, $\hat{I}_h^H$, can be used to restrict solution values if a higher order I_h^H is used for the residual, because this approximate coarse grid solution is subtracted from the update to produce an increment and is only needed for the nonlinearity of the operator (e.g., $\hat{I}_h^H = 0$ recovers CS multigrid).

3.2. **Full Multigrid Algorithm.** An effective V–cycle reduces the error by a constant fraction and is thus an iterative method, but it can be used to build a non-iterative, asymptotically exact solver that reduces the algebraic error to the order of the discretization error. FMG starts on the coarsest grid where an inexpensive accurate solve is available, prolongates the solution to the next finest level, applies a V–cycle, and continues until a desired resolution is reached. Figure 2 shows the full multigrid algorithm, with M coarse grids and α steps of the smoother before each V–cycle, in an F(α,ν1,ν2) cycle. A higher order interpolator

u =**function** FMG
  $u_0 \leftarrow 0$
  $u_0 \leftarrow FASMGV\,(L_0, u_0, f_0)$
  for k=1:M
    $u_k \leftarrow \Pi_{k-1}^{k} u_{k-1}$
    $u_k \leftarrow S^{\alpha}(L_k, u_k, f_k)$
    $u_k \leftarrow FASMGV\,(L_k, u_k, f_k)$
  return u_0

FIGURE 2. Full multigrid

between the level solves, Π_H^h, is useful for optimal efficiency of FMG and is required if I_h^H is not of sufficient order (e.g., Π_H^h must be at least linear, for cell-centered 2^{nd}-order accurate discretizations, whereas I_h^H can be constant).

One can analyze FMG with an induction hypothesis that the ratio r of the algebraic error to the discretization error is below some value and assume that the discretization error is of the form Ch^p, where p is the order of accuracy of the discretization. Further, assume that the solver on each level (e.g., one V–cycle) reduces the error by some factor Γ, which can be proven or measured experimentally, to derive the relationship between Γ and r: $\Gamma = \frac{r}{(4r+3)}$, with $p = 2$ and a refinement ratio of two. One must use a sufficiently powerful solver such that $\Gamma < 0.25$. For instance, Adams et. al. use FMG for compressible resistive magnetohydrodynamics problems; two V–cycles were required to achieve sufficient error reduction Γ [ASB10].

3.3. **Conventional distributed memory multigrid.** Domain decomposition is a natural technique for distributed memory processing of many classes of discretized PDEs, where each subdomain is placed on a processor or memory partition and the semantics of the serial algorithm are replicated. This process starts by decomposing Ω_h into P disjoint grids ${}^p\Omega_h$ such that $\Omega_h = \bigcup_{p=1}^{P} {}^p\Omega_h$. We use a rectangular array of processes of size (P_1, P_2, P_3) and thus $P = P_1P_2P_3$. We implement boundary conditions in equation (2) with ghost cells; Ω_h is enlarged by one cell in all directions to form Ω_h^{+1}. Boundary ghost cell values are set with appropriate (linear) interpolation of interior values before each operator application.

We define the number of cells on each side of a (cube) subdomain ${}^p\Omega_h$ as the integer ${}^pN_k$ on level k, again using integers for simplicity. The total number of cells in our problems is thus $n = P_1P_2P_3{}^pN_M^3$, where ${}^pN_M$ is the number of cells in each dimension on the fine grid. A conventional distributed memory full multigrid algorithm starts with a small coarse grid on a small number of processes (e.g., one process). The coarse grid is refined and split into equally sized patches, which populate more processes. This process continues until all processes are

used, forming an octree in 3D. We continue with simple refinement once all processes are used; however, more complex distributed memory models are common [ABS+14].

4. Segmental refinement

This section describes a cell-centered segmental refinement (SR) data model or method. Segmental refinement begins with a conventional distributed memory FAS-FMG method, which is used as a "coarse" grid solver. The finest level of this solver is the "transition" level and is given a grid index $k = 0$; coarser grids have negative indices. The subsequent K fine grids are defined as SR grids. Figure 3 shows a 1D example, with two SR levels and four processes.

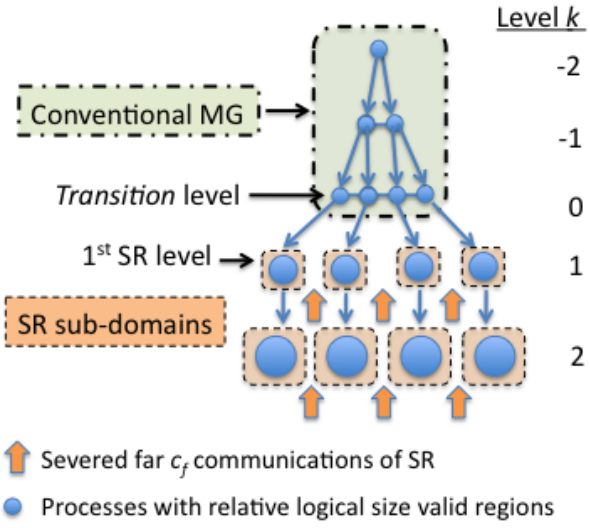

Figure 3. 1D SR data model

Non-ghost cells are defined as genuine cells and the *genuine region* is defined as ${}^{p}\Omega_k^V \equiv {}^{p}\Omega_k$. SR adds buffer cells by growing each local subdomain grid by $2J_k$ cells in each dimension. Following Brandt and Diskin, we define the length, in each dimension, of the SR buffer region J_k to be

$$J_k = J(k) = 2 \cdot \left\lfloor \frac{A + B \cdot (K - k)}{2} \right\rfloor, \tag{5}$$

where A is a constant term and B is a linearly increasing term on coarser grids. J_k is constrained to even integers to simplify restriction. The union of the genuine cells and the J_k buffer cells defines the *compute region* ${}^{p}\Omega_k^C \equiv {}^{p}\Omega_k^{V+J_k} \cap \Omega_k$, that is, the genuine region grown by $J_k h_k$ in all directions and clipped by the domain. The length of the compute region is generally ${}^{p}N_k^C = {}^{p}N_k^V + 2J_k$, where ${}^{p}N_k^V$ is the length of the compute region in grid k.

c

We define *process ghost cells* by ${}^{p}\Omega_k^G \equiv {}^{p}\Omega_k^{C+1} \setminus {}^{p}\Omega_k^C$ and subdivide ${}^{p}\Omega_k^G$ into two sets: ${}^{p}\Omega_k^{GBC} \equiv {}^{p}\Omega_k^G \setminus \Omega_k$ and ${}^{p}\Omega_k^{GSR} \equiv {}^{p}\Omega_k^G \cap \Omega_k$. The ${}^{p}\Omega_k^{GBC}$ cell values are computed with the conventional boundary condition algorithm. ${}^{p}\Omega_k^{GSR}$ cells are set during the I_0^1 prolongation process and are "frozen", in that they are not updated with the neighbor exchanges, during the rest of the multigrid process. This "freezing" is a consequence of the elided communication of SR; ${}^{p}\Omega_k^{GSR}$ cells are set with prolongation only. Define the *support of the compute region*, on grid k , of grid $k+1$ as ${}^{p}\Omega_k^F \equiv {}^{p}\Omega_{k+1}^C$; this is the region updated with the simple averaging restriction operator. Figure 4 shows a 1D example at the edge of the domain with two processes and two SR levels with the range of prolongation for one process.

The τ correction is modified to accommodate the lack of an update in the region ${}^{p}\Omega_k^C \setminus {}^{p}\Omega_k^F$. The range of prolongation is ${}^{p}\Omega^C \cup {}^{p}\Omega^{GSR}$. Figures 5 and 6 show the SR FAS-FMG algorithm with annotations for the domain of each operation.

5. Experimental observation of parameter requirements

u =**function** $FASFMGSR$
  $u_0 \leftarrow FMG$
  for k=1:K
    $u_k \leftarrow \Pi_{k-1}^{k} u_{k-1}$ $\quad {}^{p}\Omega_k^C \cup {}^{p}\Omega_k^{GSR}$
    $u_k \leftarrow S^{\alpha}(L_k, u_k, f_k)$ $\quad {}^{p}\Omega_k^C$
    $u_k \leftarrow FASMGVSR(L_k, u_k, f_k)$
  return u_K

FIGURE 5. FMG segmental refinement

u =**function** $FASMGVSR(L_k, u_k, r_k)$
  $u_k \leftarrow S^{\nu 1}(L_k, u_k, r_k)$, $\quad {}^{p}\Omega_k^C$
  $u_{k-1} \leftarrow \hat{I}_k^{k-1}(u_k)$ $\quad {}^{p}\Omega_k^F$
  $t_{k-1} \leftarrow u_{k-1}$
  $r_{k-1} \leftarrow I_k^{k-1}(r_k - L_k u_k) + L_{k-1} u_{k-1}$, $\quad {}^{p}\Omega_k^F$
  $r_{k-1} \leftarrow L_{k-1} u_{k-1}$, $\quad {}^{p}\Omega_k^C \setminus {}^{p}\Omega_k^F$
  if $k = 1$
    $w_{k-1} \leftarrow FASMGV(L_{k-1}, u_{k-1}, r_{k-1})$
  else
    $w_{k-1} \leftarrow FASMGVSR(L_{k-1}, u_{k-1}, r_{k-1})$
  $u_k \leftarrow u_k + I_{k-1}^{k}(w_{k-1} - t_{k-1})$, $\quad {}^{p}\Omega_k^C \cup {}^{p}\Omega_k^{GSR}$
  $u_k \leftarrow S^{\nu 2}(L_k, u_k, r_k)$, $\quad {}^{p}\Omega_k^C$
  return u_k

FIGURE 6. FAS V–cycle segmental refinement

This section experimentally investigates the parameters required to maintain an acceptably accurate segmental refinement FMG solver. There are several parameters that define the SR solver: the number of SR levels K and the total number levels $M+1$; A and B of equation (5); and the length of the subdomains on the transition level ${}^{p}N_0^V$.

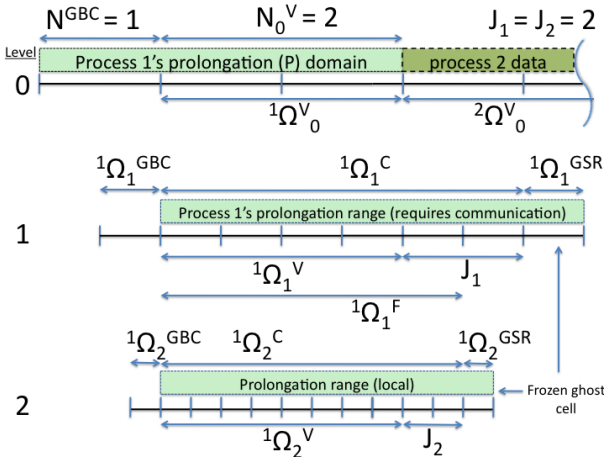

FIGURE 4. 1D 2 process SR, with the number of boundary condition cells N^{GBC}

5.1. **Model problem and solver.** We use a multigrid refinement ratio of two, piecewise constant restriction, and linear prolongation for both the FMG and V–cycle prolongation. The pre- and post-smoothers are 2^{nd}-order Chebyshev polynomials and the pre V–cycle smoother is a 1^{st}-order Chebyshev polynomial (an F(1,2,2) cycle). The solution is prescribed as $u = \prod_{i=1}^{3} (x_i^4 - R_i^2 x_i^2)$, for the Laplacian $Lu = f$, on a rectangular

domain

$$\Omega = \{x_1, x_2, x_3 \geq 0, x_1 \leq 2, x_2, x_3 \leq 1\},$$

with $R = (2, 1, 1)$ and a 4 x 2 x 2 process grid. We use a homogenous Dirichlet boundary condition and a 27-point finite volume stencil that is 2^{nd}-order accurate.

5.2. **Experiments.** We define an acceptable level of error, in the infinity norm, to be less than about 10% more than the conventional solver error (e_{conv}); the conventional solver is 2^{nd}-order convergent. We sample the parameter space of K, A, B, ${}^pN_0^V$, to find the manifold where the solver error transitions from acceptable to unacceptable. Table 1 shows the ratio (e_r) of the SR error (e_{SR}) to e_{conv} ($e_r \equiv e_{SR}/e_{conv}$) with $A = 2, 4, 6, 8$ (tables), $B = 0, 1, 2, 3$ (rows), and $\log_2 {}^pN_0^V$ and K (columns), and underlines the largest acceptable point in each column. The total number of multigrid levels can be inferred from ${}^pN_0^V$ and the process

	$\log_2 {}^pN_0^V$ (K)		
B	4(6)	3(5)	2(4)
0	17	7.2	2.7
1	2.9	2.1	1.2
2	1.5	1.2	NA
3	1.2	<u>1.1</u>	NA

(a) A=2

	$\log_2 {}^pN_0^V$ (K)		
B	4(6)	3(5)	2(4)
0	5.7	2.6	1.2
1	2.0	1.4	<u>1.0</u>
2	1.3	<u>1.1</u>	NA
3	<u>1.1</u>	1.0	NA

(b) A=4

	$\log_2 {}^pN_0^V$ (K)		
B	4(6)	3(5)	2(4)
0	2.8	1.4	<u>1.0</u>
1	1.5	<u>1.1</u>	NA
2	1.3	1.0	NA
3	<u>1.1</u>	NA	NA

(c) A=6

	$\log_2 {}^pN_0^V$ (K)		
B	4(6)	3(5)	2(4)
0	1.5	<u>1.1</u>	<u>1.0</u>
1	1.3	1.0	NA
2	<u>1.1</u>	1.0	NA
3	1.0	NA	NA

(d) A=8

TABLE 1. e_r: ratio of SR to conventional multigrid solution error

grid (i.e., $M = K + \log_2 {}^pN_0^V + \log_2 P_z = K + \log_2 {}^pN_0^V + 2$). This data shows that A and B both correlate with increased accuracy, which is expected because they both increase J. We observe that doubling the length of ${}^pN_0^V$ with K ($\log_2 {}^pN_0^V \propto K$) and increasing B with K ($B \propto K$) appears to maintain an asymptotically exact solver; we base the design of a putative asymptotically exact SR data model on this observation in §6.3.

To further investigate the effect of ${}^pN_0^V$ on error we fix $A = 8$, $B = 0$ ($J_k = 8$), and $K = 5$; the relative error as a function of ${}^pN_0^V$ is shown in Table 2. This data shows a reduction in

$\log_2 {}^pN_0^V$	5	4	3	2
$N = N_K$	1024	512	256	128
e_r	1.02	1.05	1.13	1.28

TABLE 2. $e_r({}^pN_0^V)$: $A = 8, B = 0, K = 5$

the error by a factor of about two with a doubling of ${}^pN_0^V$.

5.2.1. *Maximum segmental refinement buffer schedule.* The buffer length of the coarsest SR grid, J_1, is an important parameter because these cells require communication; these cells are the range of prolongation to the coarsest SR level and the data source for all subsequent finer grid processing. To investigate the relationship of J_1 to accuracy we test with a maximum buffer schedule (MBS), where J_1 is a parameter and rest of the SR buffers completely support grid $k = 1$: ${}^p\Omega_k^C = {}^p\Omega_k^F$ for $k < K$. This is probably not a practical buffer schedule, because the number of buffer cells increases exponentially with refinement; the MBS removes one

source of error: the lack of update of the solution and τ correction in ${}^p\Omega_k^C \setminus {}^p\Omega_k^F$. Table 3 shows the error ratio as a function of ${}^pN_0^V$ with fixed $K = 4$ and $J_1 = 4$ and the maximum buffer schedule.

$\log_2 {}^pN_0^V$	6	5	4	3	2	1
$N = N_K$	1024	512	256	128	64	32
e_r	1.02	1.05	1.11	1.25	1.4	1.9

TABLE 3. Effect of ${}^pN_0^V$ on error ratio with $J_1 = 4$, $K = 4$, with MBS

This data shows slightly less degradation of the solution with increasing ${}^pN_0^V$ than that of Table 2, but we observe a similar doubling of the error with each halving of ${}^pN_0^V$. This data indicates a dependence of accuracy on ${}^pN_0^V$, which was not recognized in the analysis of Brandt and Diskin. These are the first published multilevel numerical results of a particular segmental refinement algorithm; new algorithms and implementations should be developed with, for instance, higher order prolongation and vertex-centered discretizations to determine if this dependence can be ameliorated or if it is a fundamental property of the method.

6. Segmental refinement communication complexity

Segmental refinement inherits the computational depth of conventional distributed memory multigrid, and the coarse grid solves are identical; a more refined complexity model is required to distinguish the communication characteristics of SR from those of conventional multigrid. Mohr has analyzed the communication patterns and savings with SR and the extra computation costs for a two-level SR method [MR98, Moh00]. Brandt has presented memory complexity analysis with a $\log^D$ term for the SR buffer cells memory complexity (§8.7 [Bra84]). This section proposes a new SR data model, that we posit is asymptotically exact, and an abstract memory model that resolves the communication that is eliminated by this new SR method.

6.1. **A multigrid V–cycle communication model.** We define two types of multigrid communication: vertical inter-grid (c_V) and horizontal intra-grid (c_H) communication. A conventional distributed memory multigrid V–cycle uses 26 (c_H) messages per process in each residual, smoother, and operator application when using a 27-point stencil in 3D with a Cartesian process grid in a standard nearest-neighbor exchange process. A $V(2,2)$ cycle requires 156 horizontal messages per level, including the τ correction term, plus vertical restriction and prolongation messages in eight message phases, or bulk synchronous steps (six horizontal and two vertical). Our model focuses on these communication phases and the "distance" of each message phase.

6.2. **A memory model.** We define a "word" of data as a small patch of cells (e.g., $4^D - 32^D$ cells) and assume that each "process" computes on one data word. Consider a two-level memory model with Q words of fine grid memory, partitioned into $\sqrt{Q}$ partitions, each of size $\sqrt{Q}$. Level 1 is on-partition memory and level 2 is off-partition memory. We define near communication as communication between processes within a memory partition and far communication as communication between memory partitions.

6.3. **Proposed asymptotic segmental refinement data model.** The observations in §5 suggest that a data model that increases ${}^pN_0^V$ with the number of levels, and perhaps adds a quadratic term to equation (5), would be asymptotically exact. One could attempt to keep ${}^pN_0^V$ constant and determine an appropriate buffer schedule, but this would be less reliable and natural given our understanding of the problem. Note, a non-asymptotic model could be useful in practice because, for instance, fixing $K = 5$ reduced the size of the conventional (full communication) solver by a factor of 32K (2^{KD}), which is a significant constant.

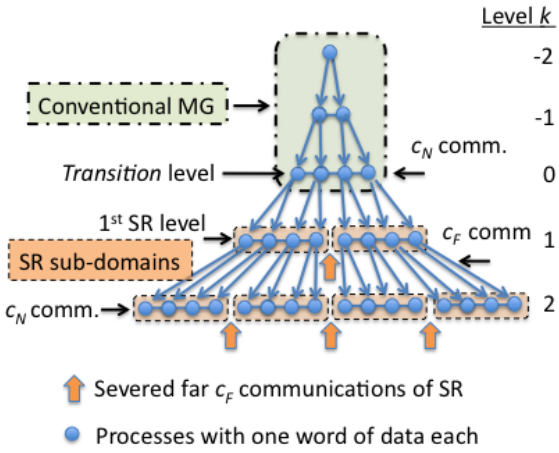

FIGURE 7. 1D example of asymptotic data model

We propose a data model that we posit would provide sufficient accuracy for an asymptotically exact solver, by extending the parallel octree of the coarse grids to the entire multigrid hierarchy, using $\sqrt{Q}$ processes in each SR subdomain and setting the size of the transition level to fit into one memory partition. With K SR levels, this model has $M = 2K$ multigrid levels and $K + 1$ conventional levels. Figure 7 shows a 1D example of this data model with two SR levels, $Q = 16$, and an SR patch length $N_0^V = 4$ where a word is one cell.

6.4. **Communication complexity.** We use the multigrid V–cycle communication model of §6.1, the machine model of §6.2, and the SR data model of §6.3 to analyze the communication complexity this segmental refinement method. We ignore FMG prolongation because there are M FMG prolongations as opposed to $\mathcal{O}(M^2)$ V–cycle restrictions and prolongations. FMG processes a V–cycle once on the finest grid, twice on the first coarse grid with $1/2^D$ as many active processes, and so on for $M+1$ levels and $(M+1)\cdot\frac{M}{2} \approx \frac{M^2}{2}$ grid visits total. This is the source of the computation depth, $log_2^2(N)$, of FMG. There are six c_H communication phases and two c_V phases per grid visit with one visit on the finest grid, two on the first coarse grid, and so on with M visits to the coarsest grid. There are about $M^2/8 = M^2/2^D$ visits on the fine (SR) half of the grid hierarchy and $3 \cdot M^2/8$ on the coarse (conventional) half of the grid hierarchy. We ignore vertical data locality and assume that all vertical communication is far communication in the finest K levels and that all communication is near communication on the coarsest $K + 1$ levels.

6.4.1. *3D Bisection bandwidth.* Briefly consider a four level memory model generated by bisecting the memory and domain of the current model. The communication complexity between these two partitions is bisection bandwidth. The highest order term of bisection bandwidth of conventional multigrid is from the ghost cell exchange on the finest grid. On a 3D cube with N cells in each dimension, $\mathcal{O}(N^2)$ cells (the area of the face between the two partitions) times the length of the ghost region ($\mathcal{O}(1)$) is communicated $\mathcal{O}(1)$ times resulting an a communication complexity of $\mathcal{O}(N^2)$.

The highest order term in the SR bisection bandwidth complexity is from the buffer region exchange on the transition level. Assume the number of buffer cells required is quadratic in K, because our data in §5 suggests this is required for a convergent solver. The "area" of data sent in this buffer cell exchange is $\mathcal{O}({}^pN_0^2) = \mathcal{O}(\sqrt{N}^2) = \mathcal{O}(N)$; it has a depth K^2 and is executed $\mathcal{O}(\log_2 N)$ times. Thus, the communication complexity is $\mathcal{O}(N \cdot K^2)\mathcal{O}(\log_2 N) = \mathcal{O}(N \cdot \log_2^3 N)$. Segmental refinement reduces the bisection communication requirements from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log_2^3 N)$.

6.4.2. *Near and far communication complexity.* Table 4 tabulates the communication complexity of conventional and SR multigrid with $M+1$ levels. The coarsest $K+1$ levels of both solvers use the same FMG solver on one memory partition. There are six c_H communication phases and two c_V phases per grid visit. The removal of far horizontal communication

Communication type	Near	Far
Coarse grids	$3 \cdot (6c_H + 2c_V)$	0
Conventional fine grids	$6c_H$	$6c_H + 2c_V$
SR fine grids	$6c_H$	$2c_V$

TABLE 4. Communication phases ($\times \log_2^2 N/8$) of conventional distributed memory multigrid and segmental refinement multigrid

complexity is the distinguishing characteristic of segmental refinement.

A given segmental refinement data model removes horizontal communication at some level of the memory hierarchy, the $6c_H$ term in far communication in Table 4 and at the arrows Figures 3 and 7. Communication, in some memory model, is used only for the vertical operators restriction and prolongation, which have tree-like graphs. Tree algorithms are efficient for the global communication required for the solve of an elliptic system. The critical observation of segmental refinement is that horizontal communication of traditional parallel multigrid is used for local processes and is not global, hence "far" communication is potentially not necessary.

We have investigated two segmental refinement data models that are "two level" in that there is one "transition" level between a conventional coarse grid solver and decoupled finer grids. One can, in principle, compose these two models, by using the method in this section as the coarse grid solver for the method in §4, and create a three level method. We speculate that one could generate an asymptotically exact "multilevel" segmental refinement method that starts with the method in this section as a "coarse grid" solver and reduces the size of ${}^pN_k^V$, by a factor of two on each finer level, resulting in just one process per SR subdomain on a fine level, thereby recovering more parallelism, and continue with the method in §4 on each process. This is a subject for future work.

7. Timing studies

This section presents scalability data on the problem in §5.1 on the Cray XC30 at NERSC, with up to 64K cores. We use 8 of the 12 cores on each socket and thus utilize 96K cores at scale, or about 75% of the machine and investigate weak scaling with 128^3 and 32^3 cells per core on the fine grid, with four and three SR levels respectively (and ${}^pN_0^V = 8$ and 4 respectively). The solver is preloaded with one solve, which verifies accuracy, followed by 8

timed solves for the 128^3 cells per core case and 512 solves for the 32^3 cells per core case to normalize times.

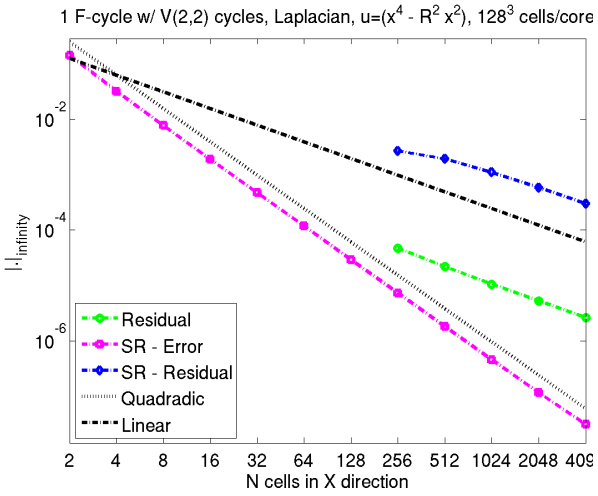

FIGURE 8. Convergence verification

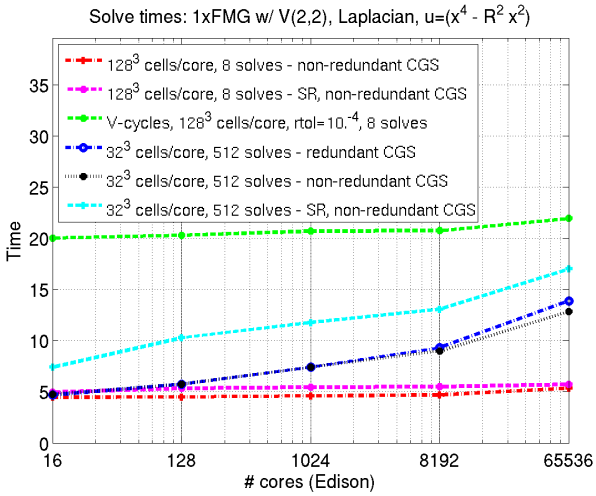
Solve times: 1xFMG w/ V(2,2), Laplacian, u=(x^4 - R^2 x^2)
128^3 cells/core, 8 solves - non-redundant CGS
128^3 cells/core, 8 solves - SR, non-redundant CGS
V-cycles, 128^3 cells/core, rtol=10.^-4, 8 solves
32^3 cells/core, 512 solves - redundant CGS
32^3 cells/core, 512 solves - non-redundant CGS
32^3 cells/core, 512 solves - SR, non-redundant CGS
Time
0 5 10 15 20 25 30 35
16 128 1024 8192 65536
cores (Edison)

FIGURE 9. Edison weak scaling

Figure 8 plots the infinity norm of the error and residual in the FMG solve and verifies that our solvers are asymptotically exact and that 2^{nd}-order accuracy is achieved, but only 1^{st}-order reduction is observed in the residual. The residuals for SR are larger than those of the conventional method but are still 1^{st}-order convergent. Figure 9 plots the solve times for the SR solver and the conventional multigrid solvers and shows modest gains in scalability with SR. The solve times for a V–cycle solve with a relative residual tolerance of 10^{-4} are also shown. Figure 10 demonstrates the stagnation in error reduction with a V–cycle solver, converged to a constant residual reduction, and that SR is maintaining perfect 2^{nd}-order accuracy.

Coarse grids can be computed redundantly, where all processors are active on all levels redundantly computing coarse grid corrections, or processors can be left idle on coarse grids. Redundant coarse grid solves result in a "butterfly" communication pattern and the idle processors result in a tree communication pattern. This approach has the advantage of requiring no communication in the prolongation phase, hence reducing the number of bulk synchronous communication steps, at the expense of sending more data with more messages overall. We observe in Figure 9 that redundant coarse grid solves are slightly slower, which suggests that larger number of messages cost more than the savings in the number of bulk synchronous phases. This could be due to contention in the network during restriction, however, the differences are small and only observable on the largest run.

8. Conclusions

We continue the work of Brandt and Diskin [BD94], with the first published multilevel numerical results of the segmental refinement multigrid method. We demonstrate that SR can maintain the semantics of textbook efficient multigrid FMG-FAS with processing that is more attractive on modern memory-centric architectures than conventional distributed

memory multigrid by decoupling fine grid processing, which improves data locality; amortizes latency costs; and reduces data dependencies. We have experimentally investigated the asymptotic behavior of SR and have found an accuracy dependance not previously recognized. We analyzed the communication complexity, with a two-level memory model, of an SR data model, where we show that the method removes horizontal communication as define by the memory model. The degree to which the memory model, on which any given SR method removes communication, is a useful performance model for any given machine can be used as a metric for the potential efficacy of the method. We experimentally verify that our SR data model is an asymptotically exact solver on 64K cores of a Cray XC30 and provide timing and scaling data.

We have observed modest improvement in scaling with SR with a simple data model that supports only a few SR levels. Future work includes developing SR data models that accommodate more levels of the memory hierarchy, testing on machines with deeper memory hierarchies and fully exploiting SR's data locality with, for instance, loop fusion [WKS+12]. A vertex-centered discretization and high order I_0^1 prolongation would be of interest to better understand the asymptotic complexity of SR and corroborate the observation of the accuracy dependance on ${}^pN_0^V$. SR may be particularly sensitive to the order of prolongation because it is used to set the "frozen" ghost cells in the SR buffer region. We have investigated a model problem; further work involves extending the application of SR to more domains, such as variable coefficient and nonlinear problems and unstructured grid problems.

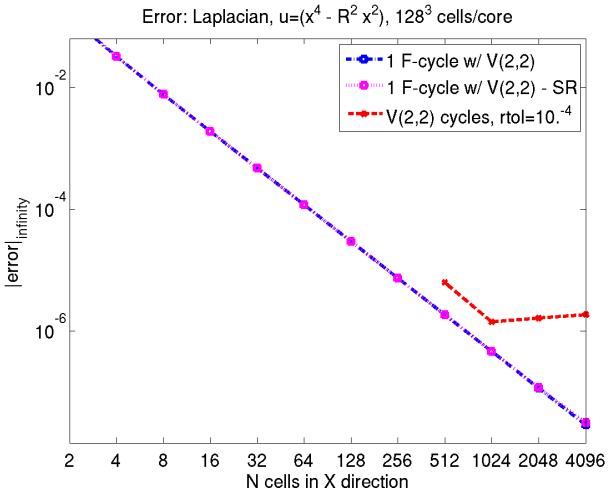

FIGURE 10. Errors of all solvers

All code, data, and run and parsing scripts used in this paper are publicly available at https://bitbucket.org/madams/srgmg.

Acknowledgments. We would like to thank Achi Brandt for his generous guidance in developing these algorithms. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and performed under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center, which is a DOE Office of Science User Facility. Authors from Lawrence Berkeley National Laboratory were supported by the U.S. Department of Energy's Advanced Scientific Computing Research Program under contract DEAC02-05CH11231.

References

[ABS+14] M. F. Adams, J. Brown, J. Shalf, B. Van Straalen, E. Strohmaier, and S. Williams, *HPGMG 1.0: A benchmark for ranking high performance computing systems*, Tech. Report LBNL-6630E, LBNL, Berkeley (https://bitbucket.org/hpgmg/hpgmg), 2014.

[ASB10] M. F. Adams, R. Samtaney, and A. Brandt, *Toward textbook multigrid efficiency for fully implicit resistive magnetohydrodynamics*, Journal of Computational Physics **229** (2010), no. 18, 6208 – 6219.

[BD81] R.E. Bank and T. Dupont, *An optimal order process for solving finite element equations*, Math. Comp. **36** (1981), 35–51.

[BD94] A. Brandt and B. Diskin, *Multigrid solvers on decomposed domains*, Domain Decomposition Methods in Science and Engineering: The Sixth International Conference on Domain Decomposition (Providence, Rhode Island), Contemporary Mathematics, vol. 157, American Mathematical Society, 1994, pp. 135–155.

[BL11] A. Brandt and O. E. Livne, *Multigrid techniques*, Society for Industrial and Applied Mathematics, 2011.

[Bra73] A. Brandt, *Multi–level adaptive technique (MLAT) for fast numerical solution to boundary value problems*, Proceedings of the Third International Conference on Numerical Methods in Fluid Mechanics (Berlin) (H. Cabannes and R. Teman, eds.), Lecture Notes in Physics, vol. 18, Springer–Verlag, 1973, pp. 82–89.

[Bra77] A. Brandt, *Multi-level adaptive solutions to boundary value problems*, Math. Comput. **31** (1977), 333–390.

[Bra84] A. Brandt, *Multigrid techniques: 1984 guide with applications to fluid dynamics*, GMD–Studien Nr. 85, Gesellschaft für Mathematik und Datenverarbeitung, St. Augustin, 1984.

[Din79] N. Dinar, *Fast methods for the numerical solution of boundary value problems*, Ph.D. thesis, Weizmann Institute of Science, Rehovot, Isreal, 1979.

[Fed61] R. P. Fedorenko, *A relaxation method for solving elliptic difference equations*, Z. Vycisl. Mat. i. Mat. Fiz. **1** (1961), 922–927, Also in U.S.S.R. Comput. Math. and Math. Phys., 1 (1962), pp. 1092–1096.

[Moh00] M. Mohr, *Low communication parallel multigrid*, Euro-Par 2000 Parallel Processing (Arndt Bode, Thomas Ludwig, Wolfgang Karl, and Roland Wismller, eds.), Lecture Notes in Computer Science, vol. 1900, Springer Berlin Heidelberg, 2000, pp. 806–814 (English).

[MR98] M. Mohr and U. Rude, *Communication reduced parallel multigrid: Analysis and experiments*, Tech. Report Technical Report No. 394, University of Augsburg, 1998.

[Sou40] R. V. Southwell, *Relaxation methods in engineering science*, Oxford University Press, Oxford, 1940.

[TDB01] J. L. Thomas, B. Diskin, and A. Brandt, *Textbook multigrid efficiency for the incompressible Navier–Stokes equations: high Reynolds number wakes and boundary layers*, Computers & Fluids **30** (2001), no. 78, 853 – 874.

[TOS01] U. Trottenberg, C. W. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, London, 2001.

[WKS+12] S. Williams, D. D. Kalamkar, A. Singh, A. M. Deshpande, B. Van Straalen, M. Smelyanskiy, A. Almgren, P. Dubey, J. Shalf, and L. Oliker, *Optimization of geometric multigrid for emerging multi- and manycore processors*, Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (Los Alamitos, CA, USA), SC '12, IEEE Computer Society Press, 2012, pp. 96:1–96:11.